

MIT Open Access Articles

*Applying the multivariate time-rescaling
theorem to neural population models*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Gerhard, Felipe, Robert Haslinger, and Gordon Pipa. "Applying the Multivariate Time-Rescaling Theorem to Neural Population Models." *Neural Computation* 23 (2011): 1452-1483. © 2011 Massachusetts Institute of Technology.

As Published: http://dx.doi.org/10.1162/NECO_a_00126

Publisher: MIT Press

Persistent URL: <http://hdl.handle.net/1721.1/66999>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Applying the Multivariate Time-Rescaling Theorem to Neural Population Models

Felipe Gerhard

felipe.gerhard@epfl.ch

Brain Mind Institute, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne EPFL, Switzerland, and Frankfurt Institute for Advanced Studies, 60438 Frankfurt am Main, Germany

Robert Haslinger

robhh@nmr.mgh.harvard.edu

Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA 02129, U.S.A., and Massachusetts Institute of Technology, Department of Brain and Cognitive Sciences, Cambridge, MA 02139, U.S.A.

Gordon Pipa

gpipa@uos.de

Institute of Cognitive Science, University of Osnabrueck, 49069 Osnabrueck, Germany; Max-Planck Institute for Brain Research, Department of Neurophysiology, 60528 Frankfurt am Main, Germany; Frankfurt Institute for Advanced Studies, 60438 Frankfurt am Main, Germany; and Massachusetts Institute of Technology, Department of Brain and Cognitive Sciences, Cambridge, MA 02139, U.S.A.

Statistical models of neural activity are integral to modern neuroscience. Recently interest has grown in modeling the spiking activity of populations of simultaneously recorded neurons to study the effects of correlations and functional connectivity on neural information processing. However, any statistical model must be validated by an appropriate goodness-of-fit test. Kolmogorov-Smirnov tests based on the time-rescaling theorem have proven to be useful for evaluating point-process-based statistical models of single-neuron spike trains. Here we discuss the extension of the time-rescaling theorem to the multivariate (neural population) case. We show that even in the presence of strong correlations between spike trains, models that neglect couplings between neurons can be erroneously passed by the univariate time-rescaling test. We present the multivariate version of the time-rescaling theorem and provide a practical step-by-step procedure for applying it to testing the sufficiency of neural population models. Using several simple analytically tractable models and more complex simulated and real data sets, we demonstrate that important features of the population activity can be detected only using the multivariate extension of the test.

1 Introduction

The strongly recurrent topologies of neuronal networks couple the activity of their component neurons together. Such coupling often manifests in experimentally observable correlations between the spiking of individual neurons. However, the importance of these correlations, and neuronal coupling in general, for information processing is still strongly debated. A possible first step toward understanding the computational role of neuronal correlations is to define and fit statistical models of neural population spiking that include the effect of possible mutual dependencies among neurons. After effectively characterizing the population activity through such a model, one can then proceed to ask whether these dependencies are in fact important for information processing or neural computation, or both. Numerous models and techniques for fitting them exist; however, all models must be checked for goodness of fit. In the case of neural population models, it is necessary for the activity of the population, not merely the activity of the individual neurons, to be accurately described. This is important because even if the statistics of the individual neurons are accurately described, this does not necessarily imply that the collective activity of the population is. The goal of this letter is to present a practical statistical test applicable to population models of neural spiking.

A goodness-of-fit measure must be geared toward both the model being tested and the data to which the model is being fit. In the case of population spiking activity, the data may be effectively represented as a set of point processes (binarized spike trains). Two main modeling approaches for such data have recently become popular in the neuroscience community. In the first, one directly models the joint distribution of spike patterns across neurons. An example of this approach is the pairwise Ising (maximum entropy) model first proposed for the analysis of population activity by Martignon et al. (2000) and others (Schneidman, Berry, Segev, & Bialek, 2006). Using this technique, they were able to demonstrate that significant pairwise correlations are present in at least some neuronal populations. Ising models treat the neuronal population activity as stationary in time (Tang et al., 2008), what is modeled is a time-independent joint probability distribution of the many-neuron spiking pattern. However, neural population activity is dynamic and highly variable with time, and a classical Ising model cannot capture this.¹

¹It should be noted that Marre and others have extended the Ising model to include spike history dependence by adding spikes in time bins up to some truncated past to the Ising model and effectively inflating the population size (Martignon et al., 2000; Tang et al., 2008; Marre, Boustani, Frégnac, & Destexhe, 2009). This results in greatly increased model fitting times because for models of the full joint distribution, a normalization constant (partition function) must be calculated that involves a sum over the probability of all 2^K possible patterns (across K neurons) of spikes. This is extremely computationally

A second approach, which is capable of modeling dynamic changes in spiking probability at arbitrarily fine temporal scales, is the point-process-based generalized linear model (GLM) technique promoted independently by Brown, Paninski, Pillow, and others (Brown, Nguyen, Frank, Wilson, & Solo, 2001; Brown, Barbieri, Eden, & Frank, 2003; Paninski, 2003; Kass, Ventura, & Brown, 2005; Truccolo, Eden, Fellows, Donoghue, & Brown, 2005; Pillow, 2007). This method models the time-varying spiking probability (conditional intensity function) of each neuron in the population independently but conditioned on the past history of all spikes in the neuronal population. The assumption being made is one of causality: due to the time delays involved in the synaptic transmission of information, neurons interact with some slight time lag.² Thus, although the model is independently defined and fit to the spikes of each neuron, the collection of models describing all neurons still constitutes a population model. This approach has also been used to characterize pairwise dependencies in neuronal populations and demonstrate their importance for describing the spike statistics of individual neurons (see, e.g., Pillow et al., 2008). But what about the importance of the correlations for the collective spike statistics of the population?

This is the issue addressed by this letter: How can one determine how accurately a set of single-neuron models describes the population statistics? We describe and present a practical methodology for such a statistical test, based on a multivariate form of the time-rescaling theorem for point processes. In its univariate form, the time-rescaling theorem states that if the conditional intensity function of a point process is known, then its interspike intervals (ISIs) can be transformed, or rescaled, as an integral over that conditional intensity function so that they are independent and exponentially distributed. In contrast, the multivariate form of the theorem states that the ISIs of a population of point processes can be rescaled so that not only are each neuron's rescaled ISIs exponentially distributed and independent of each other but they are also independent of all other neurons' rescaled ISIs.

We propose a step-by-step procedure for turning the multivariate version of the time-rescaling theorem into a goodness-of-fit measure for point-process-based population models. We demonstrate that even if each neuron's own spiking activity is sufficiently described by its conditional intensity function, the collective activity of the population may not be and that the distinction between the two cases can be made by applying the multivariate test. Toward this end, we apply the tests derived from both the univariate and multivariate time-rescaling theorem to simplified toy models, more complicated many-neuron models, and also spikes from real

intensive and must usually be done using Monte Carlo techniques, although we note that several mean field methods have recently been proposed (Broderick, Dudik, Tkacik, Schapire, & Bialek, 2007; Roudi, Tyrcha, & Hertz, 2009).

²The size of this lag need be no more than one time bin, and the bin size may be arbitrarily small.

neuronal populations recorded in awake monkey V1 during natural scenes stimulation. Thus, the multivariate extension of the time-rescaling theorem provides a means for more accurately detecting situations in which couplings between neurons significantly modulate the population dynamics and are likely to be of importance for neuronal information processing.

2 Methods

2.1 Point-Process-Based Neural Models. Data from neurophysiological experiments are often gathered in the form of a sequence of recorded action potentials. It is common to ignore the exact temporal profile of the action potentials and treat them as discrete events localized in time, or in mathematical terms, as a point process. One technique for quantifying the statistics of a point process is to describe it in terms of its conditional intensity function $\lambda(t|H_t)$ (Daley & Vere-Jones, 2002). This generalizes the concept of an instantaneous firing probability (rate) by conditioning on the past history of the point process H_t and may be formally written as

$$\lambda(t|H_t) = \lim_{\Delta \rightarrow 0} \frac{P[N(t + \Delta) - N(t) = 1|H_t]}{\Delta}, \quad (2.1)$$

where $N(t)$ is the counting function of the point process, stating the number of events inside the time interval $[0, t)$. If known exactly, the conditional intensity completely characterizes the statistical properties of the individual point process under weak assumptions. In practice, one generally hypothesizes a parametric form (which may be a function of stimuli, the neurons' own spiking history, and so on) and fits it to the recorded spikes, generally using maximum likelihood estimation.³ This approach has become standard for simultaneously quantifying the influence of multiple external covariates (such as stimuli) and also the neuron's own spike history on the output spikes (Brillinger, 1988; Chornoboy, Schramm, & Karr, 1988; Kass & Ventura, 2001; Brown et al., 2003; Paninski, 2003; Truccolo et al., 2005; Okatan, Wilson, & Brown, 2005; Kass et al., 2005; Rigat, de Gunst, & van Pelt, 2006; Pillow, 2007; Stevenson, Rebecco, Miller,

³Fitting the parameters of the conditional intensity function can be accomplished by maximizing the data likelihood, which may be expressed in terms of the conditional intensity function:

$$\log L = \sum_{t_n} \log \lambda(t_n|H_{t_n}) - \int_0^T \lambda(t|H_t) dt,$$

where $\{t_n\}$ is the set of spike times in the interval $[0, T]$.

& Körding, 2008; Pillow et al., 2008; Czanner et al., 2008; Wojcik, Mochol, Jakuczán, Wypych, & Waleszczyk, 2009).

Multiple spike trains from a neural population can be treated as a set of point processes, where each individual point process can potentially depend on the others. For such neuronal population activity, the aim is to find a statistically sufficient description of the collective activity of the population, not simply the individual spike trains of each neuron. One obvious extension is to account for mutual dependencies between neurons by conditioning each neuron's spiking on the past spiking history $\{H_t\}$ of all neurons. Formally, the conditional intensity function of each neuron in the population is now written as

$$\lambda_i(t|\{H_t\}) = \lim_{\Delta \rightarrow 0} \frac{P[N_i(t + \Delta) - N_i(t) = 1|\{H_t\}]}{\Delta}. \quad (2.2)$$

The activity of the neural population is described by the set of conditional intensity functions, one for each of the K neurons in the population:

$$\{\lambda_1(t|\{H_t\}), \lambda_2(t|\{H_t\}), \dots, \lambda_K(t|\{H_t\})\}. \quad (2.3)$$

Models that describe activity of only a single neuron are referred to as univariate models. In contrast, models that attempt to explain the firing of a whole population of neurons are referred to as multivariate models. This distinction is irrespective of whether the models make use of only the neuron's own spiking activity H_t or that of other neurons or external variables.

2.2 The Univariate Time-Rescaling Theorem. Regardless of whether a univariate or multivariate model is being considered, it must be tested for compatibility with the observed data. One possible test applicable to point process models, such as spike trains, is based on the time-rescaling theorem. This theorem and the associated test were introduced by Brown et al. to the neuroscientific community to evaluate univariate models, that is, models of single-neuron spike trains (Brown, Barbieri, Ventura, Kass, & Frank, 2002).⁴ The theorem is a special case of the theorems discussed by Meyer (1971) and Papangelou (1972) and has been applied to point process models in other scientific areas since then, for example, seismology (Ogata, 1988). We now briefly review the univariate version of the time-rescaling theorem before presenting the multivariate version and describing a new procedure for using the theorem to test multivariate models.

⁴It appeared as a "rate-rescaling theorem" in Barbieri, Quirk, Frank, Wilson, and Brown (2001).

The univariate time-rescaling theorem states that the ISIs of an orderly point process can be transformed, or rescaled, so that the rescaled process is Poisson with unit rate. That is, the rescaled ISIs are independent and exponentially distributed (Brown et al., 2002).⁵ This rescaling takes the form

$$\tau_n = \int_{t_{n-1}}^{t_n} \lambda(t|\{H_t\}) dt, \quad (2.4)$$

where $\{t_n\}$ is the set of spike times and $\lambda(t|\{H_t\})$ is the conditional intensity function of the single neuron that is to be modeled. Several proofs of this theorem exist, and we refer readers to Brown et al. (2002) for details. Although the τ_n can be compared to the exponential distribution, it is useful to note that a second transformation will rescale the ISIs to a uniform distribution:

$$z_n = 1 - e^{-\tau_n}. \quad (2.5)$$

General practice is to sort the transformed variables z_n into ascending order and plot them along the y -axis versus the uniform grid of values $b_n = \frac{n-0.5}{N}$, where N is the number of interspike intervals and $n = 1, \dots, N$. If the transformed variables z_n are indeed uniformly distributed, then the points should lie along the 45 degree line. Analytically defined confidence bounds for the maximal deviance from the diagonal can be determined and for even small sample sizes ($N > 35$) are well approximated by $b_n \pm 1.36/N^{1/2}$ for the 95% confidence level (Massey, 1951; Brown et al., 2003). Essentially the cumulative density function (CDF) of z_n s is being plotted against the CDF of the uniform distribution (the b_n s). It can be thought of as a visualization of a Kolmogorov-Smirnov (KS) test, which compares two CDFs and is usually referred to as a KS plot. In practice, the residual $z_n - b_n$ is plotted against b_n . For this differential KS plot, the points should be close to the x -axis inside the horizontal confidence bounds $\pm 1.36/N^{1/2}$. (See Figure 1B for an example of such a plot.)

Beyond being uniformly distributed, the z_n s also have to be independent of each other. Since the independence of the quantities of a time series is hard to estimate, analysis is usually constrained to look for dependencies between subsequent intervals via a scatter plot or inspecting the autocorrelation structure of interspike interval lengths (Truccolo et al., 2005). For

⁵The time-rescaling theorem applies exactly when continuous time is used and the spikes are defined as instantaneous events. As a practical matter, most statistical models discretize time into bins, and this discretization can cause a well-fitted model to, in certain instances, be erroneously rejected if the firing rate is too high compared to the bin size. Two of the authors have recently proposed a discrete-time version of the time-rescaling theorem that eliminates such problems (Haslinger, Pipa, & Brown, 2010).

a scatter plot, z_{n+1} is plotted against z_n , and in case of independence, the points should uniformly fill the unit area.

Any of these tests is designed to detect a specific departure of the rescaled times from the Poisson structure. Passing all tests is no guarantee that the rescaled times are indeed Poissonian. For example, noting that the (linear) autocorrelation between intervals up to a certain lag is not significantly different from zero implies neither uncorrelatedness for arbitrary lags nor independence, since uncorrelatedness is a weaker condition than independence; hence, intervals could be uncorrelated but still dependent (see the recent developments on copula models by Onken, Grünewälder, Munk, & Obermayer, 2009, as an example of neuroscientific relevance).

The KS test that is used to test the exponentiality of the rescaled intervals is adapted from the original work of Brown et al. (2002). However, we note that the KS test is not the only one suitable for comparing a set of samples with a proposed theoretical distribution (for a review, see Cox, 1955; Cox & Lewis, 1966; Naus, 1979). One possibility is the χ^2 test, which, however, requires more samples and is less powerful than the KS test (Lilliefors, 1969; Massey, 1951). An alternative would be to plot confidence bounds based on the probability of individual data points to deviate from the diagonal, rather than testing the maximal deviance as in the KS test. More alternatives have been formulated, for example, by using the Anderson-Darling (Anderson & Darling, 1952), Kuiper (1962), or Cramer-von Mises statistics (Anderson, 1962). One may ask why there is such a variety of different tests. The answer is simple: There are many alternative hypotheses, against the null hypothesis, which states that the samples come from a given distribution. Generally, possible deviances are even more extensive for the hypothesis of a Poisson process. Each test is particularly sensitive to certain alternative hypotheses. A goodness-of-fit procedure should reject proposed models across a broad range of possible deviances from the true model. For this purpose, the KS statistics have been proven to be robust. If, however, the expected shape of the deviance is known, other tests than the KS test might be more powerful (Daley & Vere-Jones, 2002).

The time-rescaling transformation is turned into a goodness-of-fit measure by designing statistical tests that allow the rejection of the null hypothesis that the data from the single neuron were generated by the proposed model. However, if we are recording from multiple neurons, a more general version of the theorem applies, and more rigorous tests can be performed.

2.3 The Multivariate Time-Rescaling Theorem. Within the conditional intensity function formalism, modeling the activity of several neurons at the same time amounts to modeling the conditional intensity functions of each neuron. A conceivable strategy to test such a population model would be to apply the univariate time-rescaling theorem to each of the neurons separately and reject the model if any of the individual tests fails (Truccolo et al., 2005; Rigat et al., 2006). However, this strategy may be

insufficient since any test of the collective statistics must take into account the correlations across spike trains. A simple example serves to make this clear. Consider a unit rate Poissonian spike train and a second spike train that is an exact copy of the first one, except that all spikes are shifted in time by a constant lag. Each of the spike trains can obviously be individually described as a unit rate Poisson process, and such a population model would pass the univariate version of the time-rescaling test (apart from false rejections that are controlled by the significance level of the test). This description strikingly misses the strong temporally lagged correlation between the two spike trains. The population model is, in short, wrong. A refined version of this example is more thoroughly discussed in section 3.1.

Without a test that depends on the collective statistics, it is impossible to determine if all interactions are correctly described, whether or not the entire population is included in the conditioning. Since the univariate time-rescaling theorem is not sufficient to test population models, it needs to be generalized. How does the time-rescaling theorem extend to the multivariate case? As stated in the original work of Meyer (1971) and Papangelou (1972), each rescaled point process in the population should be a unit-rate Poisson train that is independent of all of the other rescaled point processes.⁶ The independence between the rescaled processes of different neurons is a nontrivial outcome of the generalized theorems. To our knowledge, this generalized theorem has yet to be applied to neuroscientific data.

It is important to distinguish between the statement of the mathematical theorem and the statistical tests that make use of it to validate neural models. Here, assuming that we have rescaled each individual spike train with its estimated population-dependent $\lambda_i(t|\{H_i\})$ according to the theorem, how do we test that these rescaled processes are independent Poisson processes, and how do we design a statistical test that is able to reject the null hypothesis that the population's collective activity is generated according to a specific model? One strategy is to note that if each rescaled spike train constitutes a Poisson process independent of the others, then their superposition should also be a Poisson process. Along this line, we now develop the testing scheme that is based on the mathematical formulation of the time-rescaling theorem and that is shown schematically in Figure 1 and summarized in section 2.4.

First, each of the K spike trains is rescaled according to equation 2.4, that is, the transformed spike times of neuron i are given by

$$t'_{i,n} = \int_0^{t_{i,n}} \lambda_i(t|\{H_i\}) dt. \quad (2.6)$$

⁶A more comprehensive proof of the same theorem has been given in Brown and Nair (1988) and Vere-Jones and Schoenberg (2004). It has also found its way into the textbook literature (Daley & Vere-Jones, 2002).

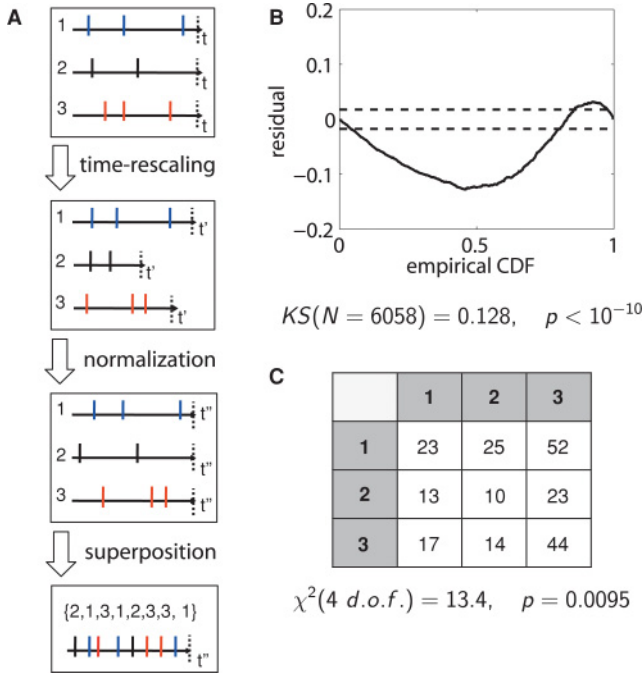


Figure 1: Procedure to test neural population models based on the multivariate time-rescaling theorem. (A) First, the individual spike trains are rescaled according to the modeled conditional intensity. The rescaled processes are normalized to equal lengths and superposed. The dotted lines indicate the end of the trial. The superposition defines a sequence of discrete marks. (B) The superposed point process can be tested against the hypothesis of a unit-rate Poisson process. The test statistic of the KS test, together (solid line) with the number of samples N , defines a p -value. This differential KS plot also displays 95% confidence intervals (dashed lines). (C) The independence of the mark sequence is tested by cross-tabulation consecutive pairs of marks and performing a χ^2 test on the observed frequencies. The test statistic χ^2 together with the degrees of freedom (d.o.f.) define the p -value for this test.

According to the theorem, all rescaled processes should be unit-rate Poissonian. Hence, the first step is to apply the univariate testing procedure to each of the rescaled processes using the methods described in section 2.2. Since more than one statistical test is performed, an appropriate multiple-testing correction has to be applied on the significance level. One common choice is to use Bonferroni-corrected⁷ significance levels for the individual tests

⁷When a statistical test is repeated K times and the global false-positive rate shall be bounded by α , the Bonferroni-corrected significance levels for each individual test

(e.g., 5/ K %). If any of these tests fails, at least one neuron is not properly modeled and the overall population model has to be rejected.

The next step is to test the independence of the rescaled processes, as required by the multivariate theorem. To this end, we note that a superposition of Poisson processes on the same domain leads again to a Poisson process with a rate equal to the sum of the individual rates. The rescaled processes, however, in general have different lengths; their domain is given by: $[0, T_i^*]$ with $T_i^* = \int_0^T \lambda_i(t|\{H_i\}) dt$, where T is the length of the trial in the original process. To obtain support on the unit domain, each transformed spike time of neuron i is multiplied by a factor of $\frac{1}{T_i^*}$. All processes can now be superposed into a single point process on the unit domain. To recover unit rate, time is finally rescaled with a factor of $\sum_{j=1}^K T_j^*$. The Poissonian nature (independent and exponentially distributed intervals) of the superposed process can be tested with the same tools that are used for the univariate version, notably the KS test on the intervals and scatter plots of intervals for a qualitative assessment of the independence of consecutive intervals.

Even if the individual processes and their superposition are Poisson processes, this does not guarantee the independence of the processes (Jacod, 1975, but see He & Wang, 1996). Therefore, as a final test, we define a mark sequence $\{m_k\}$ by placing discrete marks m_k on each spike denoting the corresponding neuron it comes from ($m_k \in \{1, 2, \dots, K\}$). If the individual point processes are Poisson and independent, the resulting superposition will also be a Poisson process with unit rate. Moreover, the marks will form an independent multinomial series with weights given by $\pi_i = \frac{T_i^*}{\sum_{j=1}^K T_j^*}$. The mark sequence $\{m_k\}$ should be an independent and identically distributed (i.i.d) sample from a multinomial series, that is, a random sequence of numbers. One out of many conceivable tests of this is based on cross-tabulating the pairs of consecutive numbers that appear in the sequence. Let $\hat{\pi}_i = \frac{N_i(T)}{\sum_j^K N_j(T)}$ denote the sample frequencies, based on the observed number of spikes. If the sequence is random, hence uncorrelated, the frequency of observing a particular pair of numbers is given by $\hat{\pi}_i \hat{\pi}_j$. Let the entries of the cross-table c_{ij} be normalized by the total number of entries. The deviation from the observed and theoretical frequencies is assessed by a χ^2 test whose test statistic is given by $\chi^2 = \sum_i^K \sum_j^K \frac{(c_{ij} - \hat{\pi}_i \hat{\pi}_j)^2}{\hat{\pi}_i \hat{\pi}_j}$ and follows a χ^2 distribution with $(K - 1)^2$ degrees of freedom. The population model is rejected if either the KS test or the χ^2 test rejects its null hypothesis based on a significantly low p-value.

are given by α/K . Note that in the context of goodness-of-fit evaluation, false positives correspond to erroneously rejecting a true model. Another possibility is to control the false discovery rate by the Benjamini-Hochberg (Benjamini & Hochberg, 1995) or Simes's procedure (Simes, 1986), for example.

2.4 Proposed Test Procedure. The statistical tests presented in section 2.3 are summarized in the following step-by-step procedure:

1. Apply the time-rescaling transformation to each spike train. Test each of the K rescaled processes for its Poissonian structure (using the classical univariate testing procedure of section 2.2).
2. Let T_i^* be the total rescaled time of each neuron i — $T_i^* = \int_0^T \lambda_i(t) \{H_t\} dt$. Denote by $\pi_i = \frac{T_i^*}{\sum_{j=1}^K T_j^*}$ the ratio of the total rescaled time of neuron i compared to the total rescaled time of the entire population. Normalize the rescaled ISIs of each neuron i by $1/\pi_i$ so that its average rescaled firing rate becomes π_i (as opposed to 1).
3. Superimpose all the rescaled processes into a single point process. Place discrete marks m_k at each spike denoting the corresponding neuron it comes from ($m_k \in \{1, 2, \dots, K\}$).
4. Test that the superimposed process is Poisson with unit rate using the same tests as were applied for the univariate case.
5. Test the independence of the mark series $\{m_k\}$ using cross-tabulation of pairs of subsequent marks and a χ^2 test.

3 Results

We now present a series of example cases for which a population of neurons passes the univariate time-rescaling theorem but not the multivariate extension. These will clearly show that unless the multivariate versions of the tests are used, neuronal models that ignore interactions may be judged adequate even when the interactions between neurons are in fact significant and strong. We begin with a simple toy model of two neurons that are mutually coupled. Two more analytically tractable examples are presented in the appendix. Next, we present results using simulated data generated from a many-neuron model with second-order couplings. Finally we show results obtained by fitting generalized linear models (GLMs) with cross-history terms to a population of V1 neurons recorded in awake macaque monkeys during natural scenes stimulation. Unless otherwise stated, the discrete time version of the univariate time-rescaling theorem (Haslinger et al., 2010) is used.

3.1 Example: Mutually Coupled Neurons. Assume a system of two model neurons, each mutually coupled with stochastic synaptic delays. That means that after neuron 1 emits a spike, neuron 2 will spike with a delay drawn from a stationary probability distribution. Because the coupling is mutual, the next spike of neuron 1 is generated after a delay to the second neuron's spike. The spike times $t_n^{(1)}$ and $t_n^{(2)}$ form a bivariate, simple point process. Formally, this corresponds to an alternating renewal process: intervals are drawn from one of two fixed distributions in an alternating




Figure 2: Mutually coupled neurons. Two neurons are mutually coupled and alternatively generate spikes with stochastic synaptic delays that are drawn from two gaussian distributions (shaded area). Example spike trains are shown with model parameters $\mu_1 = 1$ s, $\mu_2 = 5$ s, $\sigma_1 = 0.1$ s, and $\sigma_2 = 0.5$ s.

fashion. Let the two delay distributions be normal with means μ_i and variances σ_i^2 (with $\sigma_i^2 \ll \mu_i$ and truncated at zero to avoid negative intervals). Without loss of generality, we define the first spike of neuron 1 to be at $t = 0$. Figure 2 shows an example spike train. For the numerical examples, we chose $\mu_1 = 1$ s, $\mu_2 = 5$ s, $\sigma_1 = 0.02$ s, and $\sigma_2 = 1$ s and $N = 10,000$ spikes for each neuron.

If we separately analyze both neurons, we will find that each is sufficiently described by a renewal process with interspike intervals that follow a normal distribution with mean $\mu = \mu_1 + \mu_2$ and variance $\sigma^2 = \sigma_1^2 + \sigma_2^2$. Furthermore, subsequent intervals are independent as they are formed by the independently drawn delay times. An appropriate conditional intensity function for such a renewal process can be derived—for example,

$$\lambda(t|H_t) = \frac{p(t - t^*)}{1 - \int_0^{t-t^*} p(u) du}, \tag{3.1}$$

where t^* denotes the time of the last spike prior to time t and $p(t)$ denotes the probability density function (PDF) on the interspike intervals (Brown et al., 2003). In our case, $p(t)$ equals the PDF of a gaussian variable:

$$p(t) = (2\pi\sigma^2)^{-1/2} \exp\left(-\frac{(t - \mu)^2}{2\sigma^2}\right). \tag{3.2}$$

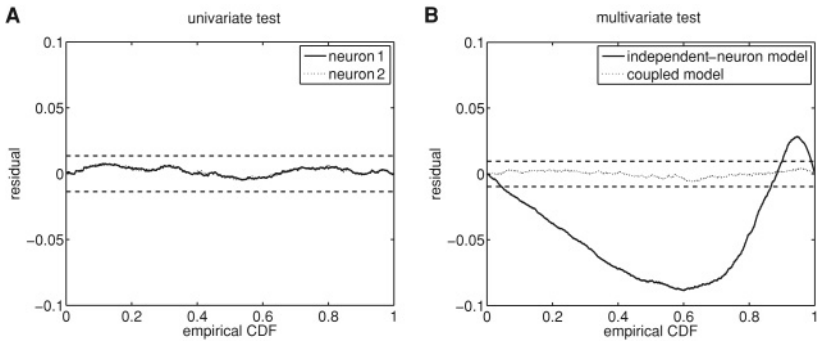


Figure 3: Mutually coupled neurons. (A) Residual KS plots for the univariate test. Both plots' (solid) lines lie within the dashed horizontal lines, which indicate 95% confidence bounds. Hence, both neurons seem to be sufficiently modeled by an independent renewal process. (B) Residual KS plots for the superimposed process of the multivariate test according to the independent-neuron assumption (solid line) and the coupled, correct model (dotted line).

For the time period of the spike trains, we can specify the conditional intensities $\lambda(t|H_t)$ according to equation 3.1. Using the univariate time-rescaling theorem of equation 2.4, a set of rescaled spike times is obtained from which the intervals τ for each of the two neurons are formed. The further transformation $z = 1 - \exp(-\tau)$ should yield uniformly distributed values in case of a correct model specification. Note that in the case of gaussian renewal processes, z can be directly calculated from the original intervals via

$$z_i^{(1,2)} = \Phi(\tau_i^{(1,2)}, \mu, \sigma^2), \quad (3.3)$$

where $\Phi(x, \mu, \sigma^2)$ is the c.d.f (cumulative density function) of a gaussian distribution with specified mean and variance. The transformation will result in rescaled spike trains that form unit-rate Poisson processes (up to statistical fluctuations). This can be numerically confirmed by the KS test (to test the exponentiality) and qualitatively using the scatter plots (to test dependence of subsequent intervals) (see Figures 3A, 4A, and 4B). From this finding, the null hypothesis of the data being generated by the univariate model cannot be rejected.

However, this model is incorrect; the spike trains of the neurons are not independent. Following section 2.4, we find that the superposed spike train is not Poissonian as a consequence of the interdependence between the two rescaled point processes. The multivariate KS test, multivariate scatter plots, and χ^2 test for independence of the mark series all fail (see Figures 3B and 4C; KS test $KS(N = 20,000) = .059$, $p < .001$, p -value




Figure 4: Mutually coupled neurons. (A, and B): Scatter plots of pairs of subsequent z -values for the univariate test do not suggest any dependence structure within neurons. (C) Scatter plot of pairs of subsequent z 's for the superposed spike train of the multivariate test for the independent-neuron model. A deviance from uniformity is visible. (D) The same scatter plot for intervals of the superposed process of the correct model. For all plots, only every tenth pair is shown. Independence is indicated by a uniformly filled unit area.

for the linear correlation between intervals $p < .001$ (Pearson's $\rho = -.05$), $\chi^2(1, N = 20,000) = 1501$, $p < .001$). Hence, following the multivariate formulation of the time-rescaling theorem, we have to reject this model, which lacks interneuron interactions.

To build a sufficient model, we have to take the interactions into account. The correct coupled model is:

- For neuron 1: The i th spike has a delay to the $(i - 1)$ th spike of neuron 2 according to a normal distribution with parameters μ_2 and σ_2^2 .
- For neuron 2: The i th spike has a delay to the i th spike of neuron 1 according to a normal distribution with parameters μ_1 and σ_1^2 .

Using this model to rescale both spike trains (using the direct transformation of equation 3.3, we find that the superposed process is indistinguishable from a Poissonian one (see Figures 3B and 4D, $KS(N = 20,000) = .0043$, $p > .05$, n.s. (not significant); p -value for the linear correlation between intervals being zero $p > .05$, n. s. (Pearson's $\rho = .0048$)). Moreover, the mark distribution is an independent realization of a binomial distribution ($\chi^2(1, N = 20,000) = .96$, $p > .05$, n. s.).

3.2 Correlated Many-Neuron Model. As a more complex example, we present the application of the proposed test to a simulated data set with many interacting neurons. Several methods have been proposed to generate spike trains with a given (pair-wise) correlational structure (Niebur, 2007; Krumin & Shoham, 2009; Macke, Berens, Ecker, Tolias, & Bethge, 2009; Gutnisky & Josic, 2010). Gutnisky and Josic (2010) describe a method to generate pair-wise correlated binary vectors by sampling from a multivariate autoregressive process whose covariance matrix is a function of

the desired correlational structure. We used this method to generate simulated spike trains of length $T = 3$ s and $N = 30$ neurons. Each of them is assigned an average firing rate of $\nu = 20$ Hz. A sequence of binary vectors is generated, representing the activation of each neuron per time bin. We set the width of the time bin to 1 ms. Exact spike times are sampled according to the procedure described by Haslinger et al. (2010). This ensures that the point process of the population is orderly and the multivariate theorem can be applied. The processes are restricted to have no significant autocorrelation and cross-correlations ρ that decay exponentially in time with a time constant of $\tau = 5$ ms. The correlations ρ are assumed to be positive and of the same peak value for all pairs of neurons. The overall strength of the introduced pair-wise correlations can be controlled by modulating ρ .

We now construct two models to describe the data. The first model assumes independent Bernoulli neurons with constant firing rate and uses a logistic regression model with a single constant term to independently fit each neuron's spikes:

$$\log \frac{\lambda_{i,t} \Delta}{1 - \lambda_{i,t} \Delta} = \beta_0. \quad (3.4)$$

This and the following logistic regression models are fitted using standard maximum likelihood techniques (McCullagh & Nelder, 1989). Such a model generally passes the univariate tests regardless of the actual strength of the correlations. Thus, it is not possible to detect the pair-wise correlational structure using the univariate test. The multivariate test, however, is sensitive to the introduced correlations and reliably rejects such an independent model. To demonstrate this, we generated toy data with varying pairwise Pearson's correlation strength, ranging from $\rho = 0$ up to 0.01. For each condition, 200 trials were generated, with the individual models fitted and validated with the multivariate test. The empirical test power is defined as the percentage of repetitions for which the multivariate test correctly rejected the independent model (rejection is based on either the KS test or the χ^2 test on the mark distribution). We expect the test power to be monotonically increasing with larger correlations ρ . As can be seen in Figure 5, the multivariate test reliably detects the departure from an independent model as soon as the average pairwise Pearson's correlations are above $\rho \approx 0.004$. The typical strength of the correlations is in agreement with typical values found in experimental recordings (Ecker et al., 2010). Schneidman et al. (2006) have shown that weak pairwise correlations that are homogeneously present in the neural population can have a strong effect on the observed global activity pattern. This might explain the sensitivity of the multivariate time-rescaling theorem and its associated tests even for rather small absolute values of correlations. Note that the evidence necessary to detect such correlations will depend on both the length of the




Figure 5: Simulated many-neuron data set. Depending on the pairwise correlation strength (Pearson’s ρ) in the data set, the test power of the multivariate time-rescaling test is shown. The test power is defined as the fraction of 200 randomly generated trials for which the multivariate theorem rejected the independent-neuron model, based on either the KS test or the χ^2 test on the mark sequence. A sigmoidal function with a baseline of 0.05 is fitted to the data points. The threshold (for which the test power exceeds 50%) is around $\rho = 0.004$ (indicated by the dashed lines). The vertical bars denote 95% confidence intervals on the test power estimates.

simulated data set and the number of neurons for which the correlations are present.

In our second model for the population activity, we construct a set of logistic regression models in which the recent spiking activity of other neurons is taken into account. Here, the spiking probability per time bin t is modeled as a nonlinear transformation of a linear sum of activities of the other neurons:

$$\log \frac{\lambda_{i,t} \Delta}{1 - \lambda_{i,t} \Delta} = g_{\text{ensemble},t}. \tag{3.5}$$

The ensemble-dependent part consists of $m_e = 6$ basis splines for the history of each other neuron. Knot points are spaced on a support of up to 15 ms. Let the j th spline be of shape $A_j(\Delta t)$. Then the contributing term to the above equation is

$$g_{\text{ensemble},t} = \sum_{k,k \neq i}^K \sum_j^{m_e} \sum_{t_n^k < t} \beta_{j,k} A_j(t - t_n^k), \tag{3.6}$$

where $\{t_n^i\}$ denotes the spike times of neuron i , and β_j denotes the coefficients of the regression model. All model parameters are estimated using maximum likelihood. Such a logistic regression model completely captures the pair-wise cross-correlations present in the data. If each neuron is modeled this way, the population model is not rejected by the multivariate test for arbitrary correlation strengths (e.g., for $\rho = 0.01$: KS test $KS(N = 1843) = .012, p > .05, n. s., \chi^2(841, N = 1843) = 866, p > .05, n. s.$)

3.3 Natural Stimulation of Macaque V1. Next, we demonstrate the use of the multivariate time-rescaling procedure using spikes from a population of neurons recorded simultaneously in monkey V1 during a natural stimulation paradigm. The experimental details are described in the appendix. An analysis from one recording session containing 42 neurons and 17 trials with a length of 5 seconds each (42,749 spikes total) is presented here. We fit two logistic regression models both without and with cross-couplings between neurons. The uncoupled model included a baseline firing rate, a spike history term and a stimulus dependency (Truccolo et al., 2005):

$$\log \frac{\lambda_{i,t} \Delta}{1 - \lambda_{i,t} \Delta} = \beta_0 + g_{\text{history},t} + g_{\text{stimulus},t}. \quad (3.7)$$

The history-dependent part consists of $m_h = 11$ basis splines. Knot points are spaced on a logarithmic scale up to 64 ms. Let the j th spline be of shape $A_j(\Delta t)$. Then the contributing term to equation 3.7 for neuron i is

$$g_{\text{history},t} = \sum_j^{m_h} \sum_{t_n^i < t} \beta_{h,j} A_j(t - t_n^i), \quad (3.8)$$

where $\{t_n^i\}$ denotes the spike times of neuron i . The stimulus term is modeled using basis splines so that the firing rate can vary as a function of the time since stimulus onset. $m_s = 74$ basis splines $B_j(t)$ with an equidistant spacing over the trial length were used. The stimulus term is thus given by

$$g_{\text{stimulus},t} = \sum_j^{m_s} \beta_{s,j} B_j(t). \quad (3.9)$$

$\beta_0, \beta_{h,j}$, and $\beta_{s,j}$ are parameters determined by maximum likelihood estimation using the standard iteratively reweighted least-squares algorithm (McCullagh & Nelder, 1989). All 42 independent neuron models passed univariate KS tests (significance level 5/42%). However the multivariate procedure rejected the set of independent neuron models as a sufficient model of the collective population activity (see Figure 6A; KS test $KS(N = 81085) = .0096, p < .001, \chi^2(1681, N = 81085) = 1731, p > .05, n. s.$)

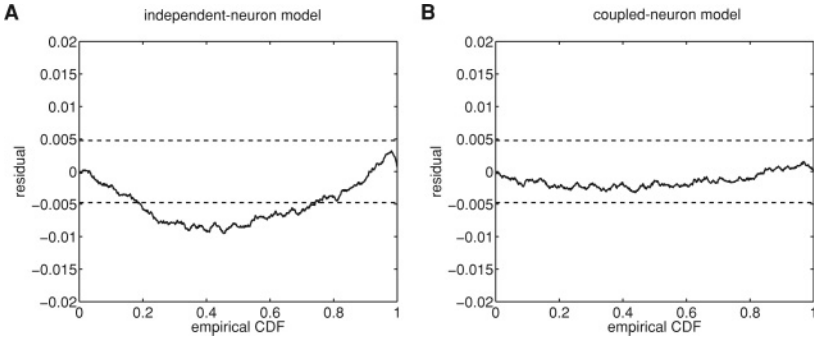


Figure 6: Monkey V1 data. (A) Residual KS plot for the superposition of the 42 spike trains using the independent neuron GLM. The multivariate test rejects this model as the line crosses the 95% confidence bounds. (B) Residual KS plot for the superposition of the 42 spike trains using a model with cross-interactions. The rescaled process lies completely within the confidence bounds.

As a second attempt at characterizing the population activity, we constructed a model that included cross-interaction terms:

$$\log \frac{\lambda_{i,t} \Delta}{1 - \lambda_{i,t} \Delta} = \beta_0 + g_{\text{history},t} + g_{\text{stimulus},t} + g_{\text{ensemble},t}. \quad (3.10)$$

The ensemble contribution is the past activity of the other recorded cells. Only interactions between cells recorded from different electrodes were included. The term $g_{\text{ensemble},t}$ is functionally similar to the self-history term but sums over all the past histories of the other neurons:

$$g_{\text{ensemble},t} = \sum_{k,k \neq i}^K \sum_j^{m_e} \sum_{t_n^k < t} \beta_{j,k} A_j(t - t_n^k). \quad (3.11)$$

In the above, the same spline filters $A_j(\Delta t)$ were used as with the self-history terms. This model with cross-couplings yields 100% performance on the univariate testing, as did the uncoupled model. This time, however, the multivariate test does not reject the population model (p -value for interspike intervals of superposed process, KS test $KS(N = 81085) = .0033$, $p > .05$, n. s., $\chi^2(1681, N = 81085) = 1675$, $p > .05$, n.s.; see Figure 6B). The cross-interaction terms thus improve the model performance considerably and would not have been detected as significant using the univariate test procedure.

4 Discussion

Simultaneous recordings from neural populations have become commonplace. Although treating neurons in a population as independent is still a popular modeling approach (Nirenberg, Carcieri, Jacobs, & Latham, 2001; Oram, Hatsopoulos, Richmond, & Donoghue, 2001; Petersen, Panzeri, & Diamond, 2001; Averbach & Lee, 2003; Jacobs et al., 2009), it has become increasingly clear that interactions between neurons matter, at least for describing the collective spike statistics, and likely for neuronal computation, although this is still hotly debated. Thus robust multivariate population models (Iyengar, 2001; Brown, Kass, & Mitra, 2004; Truccolo et al., 2005) for characterizing such activity are critical, as are appropriate statistical tests for validating these models. As yet, however, few tests of the collective statistics are being employed. When modeling neural populations, most prior studies either neglected any goodness-of-fit analysis (Chornoboy et al., 1988; Brillinger, 1988; Itskov, Curto, & Harris, 2008), or used univariate tests such as the classical time-rescaling theorem separately for each modeled spike train (Truccolo et al., 2005; Rigat et al., 2006). Using both simulated and real data, we demonstrated that this approach is insufficient because it does not test the collective statistics of the population. We presented a multivariate version of the time-rescaling theorem and proposed a test methodology suitable for determining the statistical sufficiency of neural population models. Our methods generalize time-rescaling-based goodness-of-fit tests to population models and provide one option for testing their statistical sufficiency, limited only by the finite amount of data and the statistical power of the tests.

The efficacy of our methods was demonstrated using both simulated (for which the true neuronal coupling is known) and real data sets (for which it must be inferred). The examples of mutually coupled neurons (see section 3.1) and a larger population of neurons (in section 3.2) show substantial cross-neuron dependencies that might easily be missed if each neuron is looked at only in isolation. In the appendix, we present further examples of practical interest. The example with synchronous triplets covers the case of higher-order spike patterns. Assessing the significance of higher-order spike patterns has been proven difficult in other studies (Martignon et al., 2000; Pipa, Riehle, & Grün, 2007; Pipa, Wheeler, Singer, & Nikolić, 2008). Here, a model that fails to incorporate significant higher-order interactions is likely to be rejected by tests based on the multivariate time-rescaling theorem. The second example in the appendix represents the class of common-input models (Kulkarni & Paninski, 2007; Paninski et al., 2010) or more generally latent variable models in which parts of the network that contribute to spiking are unobserved (Smith & Brown, 2003; Eden, Frank, Barbieri, Solo, & Brown, 2004; Jackson, 2004; Nykamp, 2007; Koyama & Paninski, 2009). Here, the proposed test procedure is able to detect the dependencies that are introduced by the common input or globally acting latent processes

and is therefore a strong tool for these kind of studies. Finally, we applied our method to a set of experimental data obtained from the visual cortex of the macaque monkey. It demonstrates that using the univariate theorem may erroneously indicate a good fit for independent encoding models. The lack of fit is detected by the multivariate extension and is (at least partly) corrected for by including additional cross-interaction terms in the GLM.

One has to carefully distinguish the mathematical theorem and the statistical tests that are performed. The multivariate theorem states that rescaling the individual spike trains with a proper model intensity results in a set of statistically independent Poisson processes. We chose a particular procedure to test this hypothesis by superposing all rescaled processes to obtain again a Poisson process with an independent mark sequence. Since in general, the superposition of arbitrary independent point processes will asymptotically result in a Poisson-like point process (but see Lindner, 2006), it is necessary to check that the individual processes are Poissonian (step 1 of the test procedure of section 2.4). The time rescaling requires not only the superposed but also the individual processes to be Poissonian. In addition, through this first step, the univariate version of the time-rescaling theorem is included, and the multivariate test has by definition at least the same test power. One alternative to the superposition of the processes into a single point process after time rescaling is to study the cross-correlations between pairs of the rescaled processes. If the processes are mutually independent, their cross-correlation vanishes for arbitrary lags. While this procedure is in principle feasible, it requires on the order of K^2 tests and controls only for pair-wise correlations. Our particular choice of the testing procedure has the considerable advantage that it can also detect higher-order dependencies, and the outcome is just a single point process that has to be tested for its Poisson properties. Hence, the proposed testing method inherits all the possibilities for tests that are already available for the univariate case. Although the rescaling and superposition potentially obscure temporal correlations that have not been properly modeled, we believe that the procedure works well when correlations are stationary in time and extend over the entire recording time. The additional χ^2 test on the independence of the mark sequence is a simple and well-established statistical test, although there are certainly many other statistical tests to assess the randomness of an integer sequence (for a review, see Marsaglia, 1985).

The test procedure we have proposed is derived from the multivariate time-rescaling theorem for point processes and is applicable to all population models based on the conditional intensity formalism. Modeling the conditional intensity of a point process covers a range of model classes that have been used in the neuroscientific context. Inhomogeneous Poisson processes (Olson, Gettner, Ventura, Carta, & Kass, 2000; Kass & Ventura, 2001), as well as stationary renewal processes for the interspike interval distributions—typically involving gamma (Brown et al., 2003; Maimon & Assad, 2009), log normal (Barbieri et al., 2001) or inverse gaussian distributions (Iyengar &

Liao, 1997))—can be described using the conditional intensity. Moreover, inhomogeneous Markov interval processes (IMI) are a subclass of conditional intensity models (Johnson, 1996; Kass & Ventura, 2001; Brown et al., 2002; Muller, Buesing, Schemmel, & Meier, 2007; Wojcik et al., 2009). The conditional intensity can also be nonparametrically estimated (Berry & Meister, 1998; Jacobs et al., 2009). Finally, Poisson and Bernoulli GLMs are inside the conditional intensity framework and therefore amenable to the proposed goodness-of-fit procedure. Besides modeling spike history effects, these GLMs can study the influence of external covariates, such as stimuli or the spiking activity of neighboring neurons (Harris, Csicsvari, Hirase, Dragoi, & Buzsaki, 2003; Paninski, 2003; Kass et al., 2005; Okatan et al., 2005; Ventura, Cai, & Kass, 2005; Pillow et al., 2008; Stevenson et al., 2008; Lewi, Butera, & Paninski, 2009; Toyozumi, Rad, & Paninski, 2009). Since the conditional intensity works on the level of point processes, it cannot be applied to models for continuous data (like modeling neuronal voltage dynamics or locally averaged electrical activity of a network of neurons). It can nevertheless be used to measure the influence of local field potentials to the spiking activity of single neurons (Andersen, Musallam, & Pesaran, 2004; Quiroga & Panzeri, 2009).

A distinction should be made between absolute and relative goodness-of-fit measures. For instance, the classical time-rescaling theorem was designed as an absolute goodness-of-fit test, yielding a binary decision about a specific model. Another example is the χ^2 test on the deviance of a regression model (Montgomery, Peck, & Vining, 2006). Relative goodness of fit measures compare different models and may be useful even if none of the models corresponds to the true model that generated the data. There are, in general, many measures available for assessing relative goodness of fit of neuronal models. Most commonly, the likelihood of a part of the data that was not used for fitting is evaluated (Pillow et al., 2008; Shlens et al., 2009; Itskov et al., 2008). If the experimental paradigm allows it, a model can be judged based on its ability to decode input stimuli or other conditions (Brown et al., 2004; Kass et al., 2005; Churchland, Yu, Sahani, & Shenoy, 2007; Stevenson et al., 2008). Furthermore, information-theoretic measures like the BIC (Schwarz, 1978) or AIC (Akaike, 1973) can be used (see Truccolo et al., 2005; Czanner et al., 2008). Often the power of a model to predict times of spiking is quantified (Gerstner & Naud, 2009; Truccolo, Hochberg, & Donoghue, 2010). In the case of deterministic models, their predictions can be compared to the observed data set using spike train metrics (e.g., the Γ coincidence factor—Kistler, Gerstner, & Hemmen, 1997; Jolivet et al., 2008). In the case of stochastic models, evaluations have been made using receiver-operating characteristics (Truccolo et al., 2010) and, more extensively, the (univariate) time-rescaling theorem (see Ogata, 1988; Barbieri et al., 2001; Smith & Brown, 2003; Brown et al., 2003; Rigat et al., 2006; Koyama & Kass, 2008; Wojcik et al., 2009; Shimokawa & Shinomoto, 2009). It should be noted that both the classical and the multivariate extension of

the time-rescaling theorem can be used as a relative measure by ranking models according to the absolute value of KS statistics.

Passing any absolute goodness-of-fit test does not guarantee that one has found the true model that underlies the data. It merely indicates that the observed data are consistent with a specific model within the bounds that are set by the power of the employed statistical tests. The time-rescaling procedure can be seen as a residual analysis that detects structure in the data that have not been accounted for by the model. In this way, it serves as a diagnostic tool. For example, in the first step of the multivariate procedure, each neuron is individually tested. If one or more models are rejected at this stage, the behavior of these neurons was not accurately described. If all neurons pass the individual test but the population model is rejected based on the interval distribution of the superposed process or the mark sequence, then dependencies or higher-order interactions are contained in the residual traces and indicate an incomplete population model. By performing the superposition only on a subset of neurons, the cause of the model failure may be further localized.

In this letter, we have provided one possible test of population goodness of fit. Our methodology is simple to implement and builds on well-established univariate tests of statistical sufficiency. Other techniques should be employed as well, but the multivariate time-rescaling theorem and its derived tests are a powerful addition to the statistical sufficiency test toolbox and should be widely applicable.

Appendix: Additional Examples

Here, we give two more simple examples that are analytically tractable. With these, we demonstrate the need to extend the univariate time-rescaling theorem: Even when we can estimate the conditional intensities separately for each neuron, the testing for model sufficiency has to be done in a unified way.

A.1 Example: Neurons with Synchronous Triplets. Consider three neurons, each with its own ground firing rate $\lambda^{(g)}$. On top of that, synchronous triplet events are inserted independently with a rate of $\lambda^{(t)}$. The resulting processes are again Poisson processes with rate $\lambda^{(g)} + \lambda^{(t)}$; however, these processes will be correlated through the triplet events and are hence dependent. Note that we can summarize the three processes with one marked "simple point process," where a mark S is placed for each triplet event and $\{1, 2, 3\}$ for spikes from the corresponding ground processes with rate $\lambda^{(g)}$. The observed spike train of neuron i consists of the superposition of the events with marks S and i .

For the numerical example, we used the same stationary rate of $\lambda^{(g)} = 50$ Hz for every neuron, and triplets were inserted with a rate of $\lambda^{(t)} = 10$ Hz. The simulated spike trains had a length of $T = 200$ s, using a time




Figure 7: Neurons with synchronous triplets. Three neurons exhibit synchronous triplet events (indicated by the gray boxes) on top of independent Poissonian background processes. Example spike trains were generated with a ground rate of 50 Hz and a triplet rate of 10 Hz. Only the first 200 ms are shown.

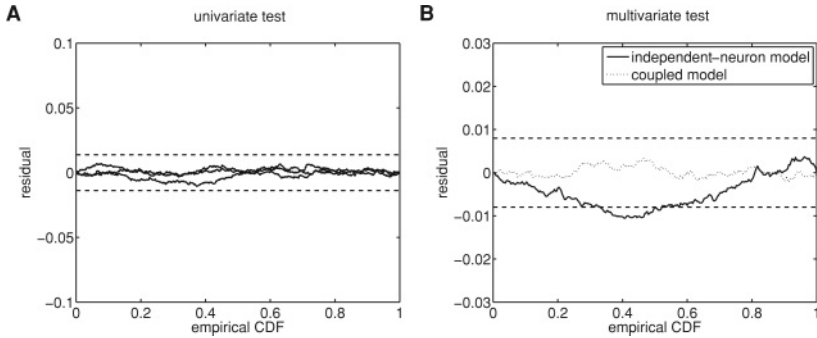


Figure 8: Neurons with synchronous triplets. (A) Residual KS plots for the univariate test. All lines lie within their 95% confidence bounds. Note that confidence bounds are dependent on the number of spikes and hence are slightly different for each neuron. (B) Residual KS plots for the superposed process of the multivariate test according to the independent-neuron assumption (solid line) and the coupled, correct model (dashed line).

discretization of $\Delta = 1$ ms. Similar to the example in section 3.2, exact spike times are sampled using the procedure of the discrete time-rescaling procedure (Haslinger et al., 2010) so that the resulting point process is simple. The synchronicity of the triplet events is therefore confined to a temporal window of $\Delta = 1$ ms. Example spike trains are shown in Figure 7.

If the univariate test was applied to each neuron individually, then a simple model such as $\lambda_i(t|H_t) = \lambda^{(s)} + \lambda^{(t)} = \text{const}$ for the activity of the individual neurons would already be statistically sufficient to describe the spikes (see Figure 8A). A constant-rate model corresponds to a uniform rescaling of all intervals with a constant factor. However, this model fails the multivariate test, as can be seen in Figure 8B ($KS(N = 29316) = .015$, $p < .001$, $\chi^2(4, N = 29316) = 44.3$, $p < .001$). Time-rescaling the individual spike trains does not affect the dependencies between the spike trains since all three spike trains are rescaled by the same constant factor.

A multivariate model that correctly describes the population can be constructed as follows. Identify the synchronous events in the population




Figure 9: Neurons with common input. $K = 6$ neurons receive synaptic input from a common Poisson process with rate 50 Hz. Neurons fire a postsynaptic spike independently from each other with probability $p = 0.2$. Only the first 200 ms of the sample are shown. The spike times of the latent background process are indicated by the shaded boxes.

recording, and model these with a homogeneous Poisson process with rate $\lambda^{(i)}$ (the events with mark S). For each spike train, consider the remaining spikes and model the occurrences of these with a Poisson model of rate $\lambda^{(g)}$ each. When time is discretized into small bins of width Δ , we get:

$$\lambda_{i,t} = \begin{cases} 1/\Delta & \text{if there was a triplet event at time bin } t. \\ \lambda^{(g)} & \text{otherwise.} \end{cases} \quad (\text{A.1})$$

After time rescaling, one obtains three processes whose Poisson nature cannot be rejected. The same holds for the superposed process ($KS(N = 29316) = .005, p > .05, \text{ n. s.}, \chi^2(4, N = 29316) = .48, p > .05, \text{ n.s.};$ see Figure 8B).

A.2 Example: Neurons with Common Input. This example is an instantiation of the multiple interaction process (MIP) of Kuhn, Aertsen, and Rotter (2003). First, a latent Poisson process with constant ground rate λ_g is generated. The spike trains of the observed neurons are independently thinned versions of λ_g ; each spike from the ground process is included in the spike train of neuron i with probability p_i . The resulting processes are Poisson processes with rate $p_i \lambda_g$. However, these processes will be correlated, hence dependent.

For the numerical example, we used a latent ground process with a stationary rate of 50 Hz and six neurons with $p_i = p = 0.2$. A possible realization of such a process is shown in Figure 9. Exact spike times are sampled using the procedure of the discrete time-rescaling procedure, and the resulting point process is simple.

A simple model for the conditional intensity such as $\lambda_i(t|H_t) = p_i \lambda_g = \text{const}$ for the activity of the individual neurons would already be sufficient as each daughter process constitutes a Poisson process itself (see Figure 10A). The time-rescaling procedure scales each spike train by a constant factor to obtain unit rate Poisson processes. This would clearly not affect the dependencies between the spike trains. Using the tests of the




Figure 10: Neurons with common input. (A) Residual KS plots for the univariate test. All lines lie within their 95% confidence bounds. (B) Residual KS plots for the superposed process of the multivariate test according to the independent-neuron assumption (solid line) and the coupled, correct model (dashed line).

multivariate time-rescaling theorem, however, we reject the independent neuron population model (see Figure 10B; $KS(N = 5897) = .13$, $p < .001$, $\chi^2(25, N = 5897) = 150$, $p < .001$).

In the correct model, the ground process is known, and the conditional intensity of the daughter processes can be modeled as

$$\lambda_i(t) = p_i \sum_j \delta(t_j), \quad (\text{A.2})$$

where $\sum_j \delta(t_j)$ is the spike train of the ground process (as a sum of delta functions at spike times t_j).

When time is discretized into small bins of width Δ , this corresponds to

$$\lambda_{i,t} = \begin{cases} p_i/\Delta & \text{if there was a spike at time bin } t \text{ of the ground process} \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.3})$$

Applying the tests based on the time-rescaling theorem, one obtains six processes whose Poisson nature cannot be rejected. The same holds for the superposed process ($KS(N = 5897) = .013$, $p > .05$, n. s., $\chi^2(25, N = 5897) = 23.1$, $p > .05$, n.s.; see Figure 10B).

A.3 Experimental Details. Data were obtained with one rhesus monkey. Experimental procedures were approved by local authorities (Regierungspraesidium Hessen, Darmstadt, Germany) and were in accord

with the guidelines of the European Community for the care and use of laboratory animals (European Union directive 86/609/EEC).

For each trial, fixation was required throughout the trial length of 5 seconds. Test stimuli consisted of natural scene movies recorded with a digital video camera (resolution 960 times 720 pixels at 30 frames per second, non-interleaved, Panasonic DVCPRO-HD format). Video sequences consisted of images of leaves, garden trees, or scenes in our laboratory obtained after a single panning movement of the camera (i.e., the video sequences always contained a single predominant global movement component).

All video clips were fully desaturated and converted into bit map image sequences cropped to a size of 936 times 702 pixels. The sequences were displayed at 100 Hz (the same frame was presented twice) using a standard graphical board controlled by ActiveStim (average luminance, $10 \text{ cd}/\text{m}^2$).

Recordings were made from the opercular region of V1 (receptive fields centers, 2.0 to 3.0 degrees eccentricity) and from the superior bank of the calcarine sulcus (10.0 to 13.0 degrees eccentricity). Electrodes were inserted independently into the cortex via guide tubes positioned above the dura (diameter, $300 \mu\text{m}$; Ehrhardt Söhne, Germany), assembled in a customized recording device. Quartz-insulated tungsten-platinum electrodes (Thomas Recording, Germany; diameter, $80 \mu\text{m}$) with impedances ranging from 0.3 to $1.0 \text{ M}\Omega$ were used to record simultaneously the extracellular activity from four to five sites in both superficial and deep layers of the cortex.

The spiking activity of small groups of neurons (MUA) was obtained by amplifying ($1000\times$) and bandpass filtering (0.7 to 6.0 kHz) the recorded signals with a customized 32 channel Plexon preamplifier connected to an HST16o25 headset (Plexon Inc., U.S.A.). Additional $10\times$ signal amplification was done by on-board amplifiers (E-series acquisition boards, National Instruments, U.S.A.). The signals were digitized and stored using a LabVIEW-based acquisition system developed in our laboratory. Spikes were detected by amplitude thresholding, which was set interactively after online visualization of the spike waveforms (typically two to three standard deviations above noise level). Spike events and corresponding waveforms were sampled at $32 \text{ kS}/\text{s}$ (spike waveform length, 1.2 ms).

Offline spike sorting was performed using a dynamic template-matching method implemented in a custom software package. Sorting was initiated by an automatic procedure that defined up to 12 different clusters. Afterward, various displays, such as tuning curves, autocorrelograms, and measurements of recording stability, were used to guide interactively which cluster to merge or delete. Only clusters well separated in 2D and 3D plots of spike principal component analysis scores were assigned to single units if a refractory period was confirmed in interspike interval distributions.

Acknowledgments

We are grateful to Sergio Neuenschwander and Bruss Lima for supplying the macaque V1 recordings discussed in section 3.3. This work was supported by NIH grant K25 NS052422-02, the Max Planck Society, and EU Grant PHOCUS, 240763, FP7-ICT-2009-C. F.G. acknowledges partial support by the Stiftung Polytechnische Gesellschaft (Frankfurt am Main, Germany) and support by the Swiss National Science Foundation under grant number 200020-117975.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Csaki (Eds.), *Second International Symposium on Information Theory* (pp. 267–281). Budapest: Akadémiai Kiado.
- Andersen, R. A., Musallam, S., & Pesaran, B. (2004). Selecting the signals for a brain-machine interface. *Current Opinion in Neurobiology*, *14*(6), 720–726.
- Anderson, T. W. (1962). On the distribution of the two-sample Cramer–von Mises criterion. *Annals of Mathematical Statistics*, *33*(3), 1148–1159.
- Anderson, T. W., & Darling, D. A. (1952). Asymptotic theory of certain goodness-of-fit criteria based on stochastic processes. *Annals of Mathematical Statistics*, *23*(2), 193–212.
- Averbeck, B. B., & Lee, D. (2003). Neural noise and movement-related codes in the macaque supplementary motor area. *Journal of Neuroscience*, *23*(20), 7630–7641.
- Barbieri, R., Quirk, M. C., Frank, L. M., Wilson, M. A., & Brown, E. N. (2001). Construction and analysis of non-Poisson stimulus-response models of neural spiking activity. *Journal of Neuroscience Methods*, *105*(1), 25–37.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B (Methodological)*, *57*(1), 289–300.
- Berry, M. J., & Meister, M. (1998). Refractoriness and neural precision. *J. Neurosci.*, *18*(6), 2200–2211.
- Brillinger, D. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, *59*(3), 189–200.
- Broderick, T., Dudik, M., Tkacik, G., Schapire, R. E., & Bialek, W. (2007). *Faster solutions of the inverse pairwise Ising problem*. Available online at <http://arxiv.org/abs/0712.2437v2>.
- Brown, E. N., Barbieri, R., Eden, U. T., & Frank, L. M. (2003). Likelihood methods for neural spike train data analysis. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach* (pp. 253–286). London: Chapman and Hall.
- Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E., & Frank, L. M. (2002). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, *14*(2), 325–346.
- Brown, E. N., Kass, R. E., & Mitra, P. P. (2004). Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nature Neuroscience*, *7*(5), 456–461.

- Brown, E. N., Nguyen, D. P., Frank, L. M., Wilson, M. A., & Solo, V. (2001). An analysis of neural receptive field plasticity by point process adaptive filtering. *Proceedings of the National Academy of Sciences*, 98(21), 12261–12266.
- Brown, T. C., & Nair, G. M. (1988). A simple proof of the multivariate random time change theorem for point processes. *Journal of Applied Probability*, 25(1), 210–214.
- Chornoboy, E., Schramm, L., & Karr, A. (1988). Maximum likelihood identification of neural point process systems. *Biological Cybernetics*, 59(4), 265–275.
- Churchland, M. M., Yu, B. M., Sahani, M., & Shenoy, K. V. (2007). Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Current Opinion in Neurobiology*, 17(5), 609–618.
- Cox, D. R. (1955). Some statistical methods connected with series of events. *Journal of the Royal Statistical Society. Series B*, 17(2), 129–164.
- Cox, D. R., & Lewis, P. A. W. (1966). *Statistical analysis of series of events*. London: Chapman and Hall.
- Czanner, G., Eden, U. T., Wirth, S., Yanike, M., Suzuki, W. A., & Brown, E. N. (2008). Analysis of between-trial and within-trial neural spiking dynamics. *Journal of Neurophysiology*, 99(5), 2672–2693.
- Daley, D. J., & Vere-Jones, D. (2002). *An introduction to the theory of point processes 1* (2nd ed.). New York: Springer.
- Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., & Tolias, A. S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965), 584–587.
- Eden, U. T., Frank, L. M., Barbieri, R., Solo, V., & Brown, E. N. (2004). Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Computation*, 16(5), 971–998.
- Gerstner, W., & Naud, R. (2009). Neuroscience. How good are neuron models? *Science*, 326(5951), 379–380.
- Gutnitsky, D. A., & Josic, K. (2010). Generation of spatiotemporally correlated spike trains and local field potentials using a multivariate autoregressive process. *Journal of Neurophysiology*, 103(5), 2912–2930.
- Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., & Buzsaki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424(6948), 552–556.
- Haslinger, R., Pipa, G., & Brown, E. (2010). Discrete time rescaling theorem: Determining goodness of fit for discrete time statistical models of neural spiking. *Neural Computation*, 22(10), 2477–2506.
- He, S., & Wang, J. (1996). Thinning of point processes, revisited. *Acta Mathematicae Applicatae Sinica*, 12(3), 278–283.
- Itskov, V., Curto, C., & Harris, K. D. (2008). Valuations for spike train prediction. *Neural Computation*, 20(3), 644–667.
- Iyengar, S. (2001). The analysis of multiple neural spike trains. In N. Balakrishnan (Ed.), *Advances in methodological and applied aspects of probability and statistics* (pp. 507–524). Boca Raton, FL: CRC Press.
- Iyengar, S., & Liao, Q. (1997). Modeling neural activity using the generalized inverse Gaussian distribution. *Biological Cybernetics*, 77(4), 289–295.
- Jackson, S. S. (2004). Including long-range dependence in integrate-and-fire models of the high interspike-interval variability of cortical neurons. *Neural Computation*, 16(10), 2125–2195.

- Jacobs, A. L., Fridman, G., Douglas, R. M., Alam, N. M., Latham, P., Prusky, G. T., et al. (2009). Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences*, 106(14), 5936–5941.
- Jacod, J. (1975). Two dependent Poisson processes whose sum is still a Poisson process. *Journal of Applied Probability*, 12(1), 170–172.
- Johnson, D. H. (1996). Point process models of single-neuron discharges. *Journal of Computational Neuroscience*, 3(4), 275–299–299.
- Jolivet, R., Kobayashi, R., Rauch, A., Naud, R., Shinomoto, S., & Gerstner, W. (2008). A benchmark test for a quantitative assessment of simple neuron models. *Journal of Neuroscience Methods*, 169(2), 417–424.
- Kass, R. E., & Ventura, V. (2001). A spike-train probability model. *Neural Computation*, 13(8), 1713–1720.
- Kass, R. E., Ventura, V., & Brown, E. N. (2005). Statistical issues in the analysis of neuronal data. *J. Neurophysiol.*, 94(1), 8–25.
- Kistler, W. M., Gerstner, W., & Hemmen (1997). Reduction of the Hodgkin-Huxley equations to a single-variable threshold model. *Neural Computation*, 9(5), 1015–1045.
- Koyama, S., & Kass, R. E. (2008). Spike train probability models for stimulus-driven leaky integrate-and-fire neurons. *Neural Computation*, 20(7), 1776–1795.
- Koyama, S., & Paninski, L. (2009). Efficient computation of the maximum a posteriori path and parameter estimation in integrate-and-fire and more general state-space models. *Journal of Computational Neuroscience*, 29, 89–205.
- Krumin, M., & Shoham, S. (2009). Generation of spike trains with controlled auto- and cross-correlation functions. *Neural Computation*, 21(6), 1642–1664.
- Kuhn, A., Aertsen, A., & Rotter, S. (2003). Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Computation*, 15(1), 67–101.
- Kuiper, N. H. (1962). Tests concerning random points on a circle. *Nederl. Akad. Wetensch. Proc. Ser. A*, 63, 38–47.
- Kulkarni, J. E., & Paninski, L. (2007). Common-input models for multiple neural spike-train data. *Network*, 18(4), 375–407.
- Lewi, J., Butera, R., & Paninski, L. (2009). Sequential optimal design of neurophysiology experiments. *Neural Computation*, 21(3), 619–687.
- Lilliefors, H. W. (1969). On the Kolmogorov-Smirnov test for the exponential distribution with mean unknown. *Journal of the American Statistical Association*, 64(325), 387–389.
- Lindner, B. (2006). Superposition of many independent spike trains is generally not a Poisson process. *Physical Review E*, 73(2), 022901.
- Macke, J. H., Berens, P., Ecker, A. S., Tolias, A. S., & Bethge, M. (2009). Generating spike trains with specified correlation coefficients. *Neural Computation*, 21(2), 397–423.
- Maimon, G., & Assad, J. A. (2009). Beyond Poisson: Increased spike-time regularity across primate parietal cortex. *Neuron*, 62(3), 426–440.
- Marre, O., Boustani, S. E., Frégnac, Y., & Destexhe, A. (2009). Prediction of spatiotemporal patterns of neural activity from pairwise correlations. *Physical Review Letters*, 102(13), 138101.
- Marsaglia, G. (1985). A current view of random number generators. In L. Billard (Ed.), *Computer science and statistics: The interface* (pp. 3–10). Amsterdam: Elsevier.

- Martignon, L., Deco, G., Laskey, K., Diamond, M., Freiwald, W., & Vaadia, E. (2000). Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Computation*, *12*(11), 2621–2653.
- Massey, F. J. (1951). The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American Statistical Association*, *46*(253), 68–78.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models*. Boca Raton, FL: Chapman & Hall/CRC.
- Meyer, P. A. (1971). Demonstration simplifiée d'un théorème de Knight. In *Séminaire de Probabilités V Université de Strasbourg* (pp. 191–195). Berlin: Springer.
- Montgomery, D. C., Peck, E. A., & Vining, G. G. (2006). *Introduction to linear regression analysis* (4th ed.). Hoboken, NJ: Wiley.
- Muller, E., Buesing, L., Schemmel, J., & Meier, K. (2007). Spike-frequency adapting neural ensembles: Beyond mean adaptation and renewal theories. *Neural Computation*, *19*(11), 2958–3010.
- Naus, J. I. (1979). An indexed bibliography of clusters, clumps and coincidences. *International Statistical Review*, *47*(1), 47–78.
- Niebur, E. (2007). Generation of synthetic spike trains with defined pairwise correlations. *Neural Computation*, *19*(7), 1720–1738.
- Nirenberg, S., Carcieri, S. M., Jacobs, A. L., & Latham, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature*, *411*(6838), 698–701.
- Nykamp, D. Q. (2007). A mathematical framework for inferring connectivity in probabilistic neuronal networks. *Mathematical Biosciences*, *205*(2), 204–251.
- Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, *83*, 9–27.
- Okatan, M., Wilson, M. A., & Brown, E. N. (2005). Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Computation*, *17*(9), 1927–1961.
- Olson, C. R., Gettner, S. N., Ventura, V., Carta, R., & Kass, R. E. (2000). Neuronal activity in macaque supplementary eye field during planning of saccades in response to pattern and spatial cues. *J. Neurophysiol.*, *84*(3), 1369–1384.
- Onken, A., Grünewälder, S., Munk, M. H., & Obermayer, K. (2009). Analyzing short-term noise dependencies of spike-counts in macaque prefrontal cortex using copulas and the flashlight transformation. *PLoS Computational Biology*, *5*(11), e1000577.
- Oram, M. W., Hatsopoulos, N. G., Richmond, B. J., & Donoghue, J. P. (2001). Excess synchrony in motor cortical neurons provides redundant direction information with that from coarse temporal measures. *Journal of Neurophysiology*, *86*(4), 1700–1716.
- Paninski, L. (2003). Estimation of entropy and mutual information. *Neural Computation*, *15*(6), 1191–1253.
- Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rahnama Rad, K., Vidne, M., et al. (2010). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, *29*(1), 107–126.
- Papangelou, F. (1972). Integrability of expected increments of point processes and a related random change of scale. *Transactions of the American Mathematical Society*, *165*, 483–506.

- Petersen, R. S., Panzeri, S., & Diamond, M. E. (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron*, 32(3), 503–514.
- Pillow, J. (2007). Likelihood-based approaches to modeling the neural code. In K. Doya, S. Ishii, A. Pouget, & R. P. N. Rao (Eds.), *Bayesian brain: Probabilistic approaches to neural coding*. Cambridge, MA: MIT Press.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207), 995–999.
- Pipa, G., Riehle, A., & Grün, S. (2007). Validation of task-related excess of spike coincidences based on NeuroXidence. *Neurocomputing*, 70(10–12), 2064–2068.
- Pipa, G., Wheeler, D. W., Singer, W., & Nikolić, D. (2008). NeuroXidence: Reliable and efficient analysis of an excess or deficiency of joint-spike events. *Journal of Computational Neuroscience*, 25(1), 64–88.
- Quiroga, R. Q., & Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nature Rev. Neuroscience*, 10(3), 173–185.
- Rigat, F., de Gunst, M., & van Pelt, J. (2006). Bayesian modelling and analysis of spatio-temporal neuronal networks. *Bayesian Analysis*, 1(4), 733–764.
- Roudi, Y., Tyrcha, J., & Hertz, J. (2009). Ising model for neural data: Model quality and approximate methods for extracting functional connectivity. *Physical Review E*, 79(5), e051915.
- Schneidman, E., Berry, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087), 1007–1012.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6(2), 461–464.
- Shimokawa, T., & Shinomoto, S. (2009). Estimating instantaneous irregularity of neuronal firing. *Neural Computation*, 21(7), 1931–1951.
- Shlens, J., Field, G. D., Gauthier, J. L., Greschner, M., Sher, A., Litke, A. M., et al. (2009). The structure of large-scale synchronized firing in primate retina. *Journal of Neuroscience*, 29(15), 5022–5031.
- Simes, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika*, 73(3), 751–754.
- Smith, A. C., & Brown, E. N. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15(5), 965–991.
- Stevenson, I. H., Rebesco, J. M., Miller, L. E., & Körding, K. P. (2008). Inferring functional connections between neurons. *Current Opinion in Neurobiology*, 18(6), 582–588.
- Tang, A., Jackson, D., Hobbs, J., Chen, W., Smith, J. L., Patel, H., et al. (2008). A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *Journal of Neuroscience*, 28(2), 505–518.
- Toyoizumi, T., Rad, K. R., & Paninski, L. (2009). Mean-field approximations for coupled populations of generalized linear model spiking neurons with Markov refractoriness. *Neural Computation*, 21(5), 1203–1243.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., & Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.*, 93(2), 1074–1089.

- Truccolo, W., Hochberg, L. R., & Donoghue, J. P. (2010). Collective dynamics in human and monkey sensorimotor cortex: Predicting single neuron spikes. *Nature Neuroscience*, *13*(1), 105–111.
- Ventura, V., Cai, C., & Kass, R. E. (2005). Trial-to-trial variability and its effect on time-varying dependency between two neurons. *J. Neurophysiol.*, *94*(4), 2928–2939.
- Vere-Jones, D., & Schoenberg, F. P. (2004). Rescaling marked point processes. *Australian and New Zealand Journal of Statistics*, *46*(1), 133–143.
- Wojcik, D. K., Mochol, G., Jakuczyn, W., Wypych, M., & Waleszczyk, W. (2009). Direct estimation of inhomogeneous Markov interval models of spike trains. *Neural Computation*, *21*(8), 2105–2113.

Received April 22, 2010; accepted October 23, 2010.